WHAT IS CLAIMED IS:

1. A system, comprising:

a plurality of nodes coupled by an inter-node network, wherein each node includes a plurality of active devices, a memory subsystem, and an address network and a data network respectively configured to convey address packets and data packets between the plurality of active devices and the memory subsystem;

wherein a memory subsystem included in a node of the plurality of nodes is configured to maintain a response indication indicating whether the memory subsystem should send a data packet corresponding to a coherency unit in response to receiving from an active device in the node an address packet requesting an access right to the coherency unit;

wherein the node is configured to store a node identifier for the coherency unit, wherein the node identifier identifies a different node of the plurality nodes in which the coherency unit is in a modified global access state.

2. The system of claim 1, wherein each node includes an interface coupled to send and receive coherency messages on the inter-node network, wherein an interface included in the node is configured to store the node identifier for the coherency unit.

3. The system of claim 2, wherein the interface is configured to store the node identifier in a global information cache that includes storage for a plurality of node identifiers for a plurality of coherency units.

4. The system of claim 2, wherein the interface included in the node is further configured to store a global access state of the coherency unit in the node.

5. The system of claim 2, wherein in response to a coherency message sent by another one of the plurality of nodes requesting an access right to the coherency unit, the interface included in the node is configured to access the node identifier and to responsively send an additional coherency message to an interface included in the different node.

6. The system of claim 5, wherein in response to receiving yet another coherency message acknowledging that the other one of the plurality of nodes gained the access right to the coherency unit, the interface included in the node is configured to update the node identifier to identify the other one of the plurality of nodes if an active device included in the other one of the plurality of nodes gained write access to the coherency unit.

7. The system of claim 6, wherein if the interface included in the node updates the node identifier, the interface is configured to send an address packet indicating a new value of the node identifier to the memory subsystem included in the node.

8. The system of claim 2, wherein the memory subsystem is configured to update the response indication in response to receiving address packets from active devices included in the node.

9. The system of claim 8, wherein the memory subsystem is configured to update the response indication to indicate that the memory subsystem should not respond to address packets requesting an access right to the coherency unit in response to receiving an address packet requesting a write access right to the coherency unit from one of the active devices included in the node.

10. The system of claim 8, wherein the memory subsystem is configured to update the response indication to indicate that the memory subsystem should respond to address

packets requesting an access right to the coherency unit in response to receiving an address packet requesting to write a new value of the coherency unit to the memory subsystem.

11. The system of claim 8, wherein the memory subsystem is configured to not update the response indication in response to address packets requesting a shared access right to the coherency unit.

12. The system of claim 8, wherein the memory subsystem is configured to send a packet corresponding to the address packet to the interface included in the node if:

the response indication indicates that the memory should not respond;

the coherency unit is in a shared global access state in the node; and

the address packet requests write access to the coherency unit.

13. The system of claim 12, wherein in response to the packet corresponding to the address packet, the interface is configured to send a coherency message requesting write access to the coherency unit to the different node identified by the node identifier.

14. The system of claim 8, wherein the memory subsystem is configured to send a packet corresponding to the address packet to the interface included in the node if:

the response indication indicates that the memory should not respond; and

the coherency unit is in an invalid global access state in the node.

15. The system of claim 1, wherein the node is configured to update the node identifier in response to receiving coherency messages from other ones of the plurality of nodes via the inter-node network.

16. The system of claim 1, wherein an active device that sends the address packet specifying the coherency unit gains the access right to the coherency unit in response to receiving the data packet from the memory subsystem if the memory subsystem sends the data packet.

17. A node for use in a multi-node computer system, the node comprising:

a plurality of client devices including a memory subsystem, an active device, and an interface configured to send and receive coherency messages on an inter-node network coupling nodes in the multi-node computer system;

an address network configured to convey address packets between the plurality of client devices;

a data network configured to convey data packets between the plurality of client devices;

wherein the memory is configured to maintain a response indication indicating whether the memory should send a data packet corresponding to a coherency unit on the data network in response to receiving an address packet requesting an access right to the coherency unit from the active device;

wherein the node is configured to store a node identifier for the coherency unit, wherein the node identifier identifies a different node in the multi-node system in which the coherency unit is in a modified global access state.

18. The node of claim 17, wherein the interface is configured to store the node identifier for the coherency unit.

19. The node of claim 18, wherein the interface is configured to store the node identifier in a global information cache that includes storage for a plurality of node identifiers for a plurality of coherency units.

20. The node of claim 18, wherein the interface included in the node is further configured to store a global access state of the coherency unit in the node.

21. The node of claim 18, wherein in response to receiving a coherency message sent by another one of the nodes requesting an access right to the coherency unit, the interface is configured to access the node identifier and to responsively send an additional coherency message to an interface included in the different node.

22. The node of claim 21, wherein in response to receiving yet another coherency message acknowledging that the other one of the plurality of nodes gained the access right to the coherency unit, the interface is configured to update the node identifier to identify the other one of the plurality of nodes if an active device included in the other one of the plurality of nodes gained write access to the coherency unit.

23. The node of claim 22, wherein if the interface updates the node identifier, the interface is configured to send an address packet indicating a new value of the node identifier to the memory subsystem.

24. The node of claim 17, wherein the memory subsystem is configured to update the response indication in response to receiving address packets from one or more active devices included in the node.

25. The node of claim 24, wherein the memory subsystem is configured to update the response indication to indicate that the memory subsystem should not respond to address packets requesting an access right to the coherency unit in response to receiving an address packet requesting a write access right to the coherency unit from one of the one or more active devices.

26. The node of claim 24, wherein the memory subsystem is configured to update the response indication to indicate that the memory subsystem should respond to address packets requesting an access right to the coherency unit in response to receiving an address packet requesting to write a new value of the coherency unit to the memory subsystem.

27. The node of claim 24, wherein the memory subsystem is configured to not update the response indication in response to address packets requesting a shared access right to the coherency unit.

28. The node of claim 24, wherein the memory subsystem is configured to send a packet corresponding to the address packet to the interface if:

the response indication indicates that the memory should not respond;

the coherency unit is in a shared global access state in the node; and

the address packet requests write access to the coherency unit.

29. The node of claim 28, wherein in response to the packet corresponding to the address packet, the interface is configured to send a coherency message requesting write access to the coherency unit to the different node identified by the node identifier.

30. The node of claim 24, wherein the memory subsystem is configured to send a packet corresponding to the address packet to the interface if:

the response indication indicates that the memory should not respond; and

the coherency unit is in an invalid global access state in the node.

31. The node of claim 17, wherein the node is configured to update the node identifier in response to the interface receiving coherency messages from other nodes in the multi-node system.

32. The node of claim 17, wherein the active device that sends the address packet specifying the coherency unit gains the access right to the coherency unit in response to receiving the data packet from the memory subsystem if the memory subsystem sends the data packet.

33. A method of operating a multi-node computer system, wherein the multi-node computer system includes a plurality of nodes coupled by an inter-node network, wherein each node includes an active device, a memory subsystem, and an address network coupling the active device and the memory subsystem, the method comprising:

a memory subsystem included in a node of the plurality of nodes receiving from an active device included in the node an address packet requesting an access right to a coherency unit;

in response to said receiving, the memory subsystem sending a responsive data packet to the active device dependent on response indication associated with the coherency unit;

the node sending a coherency message requesting the access right to a different node of the plurality of nodes in response to a node identifier identifying the different node as a node in which the coherency unit is in a modified global access state.

34. The method of claim 33, further comprising an interface included in the node storing the node identifier for the coherency unit and sending the coherency message to the different node.

35. The method of claim 34, wherein said storing comprises the interface storing the node identifier in a global information cache that includes storage for a plurality of node identifiers for a plurality of coherency units.

36. The method of claim 34, further comprising the interface storing a global access state of the coherency unit in the node.

37. The method of claim 34, further comprising:

in response to receiving a coherency message sent by another one of the nodes requesting an access right to the coherency unit, the interface sending an additional coherency message to an interface included in the different node identified by the node identifier.

38. The method of claim 37, further comprising:

in response to receiving yet another coherency message acknowledging that the other one of the plurality of nodes gained the access right to the coherency unit, the interface updating the node identifier to identify the other one of the plurality of nodes if an active device included in the other one of the plurality of nodes gained write access to the coherency unit.

39. The method of claim 38, further comprising the interface sending an address packet indicating a new value of the node identifier to the memory subsystem if the interface updates the node identifier.

40. The method of claim 34, further comprising the memory subsystem updating the response indication in response to receiving address packets from one or more active devices included in the node.

41. The method of claim 40, wherein said updating comprises updating the response indication to indicate that the memory subsystem should not respond to address packets requesting an access right to the coherency unit in response to receiving an address packet requesting a write access right to the coherency unit from one of the one or more active devices.

42. The method of claim 40, wherein said updating comprises updating the response indication to indicate that the memory subsystem should respond to address packets requesting an access right to the coherency unit in response to receiving an address packet requesting to write a new value of the coherency unit to the memory subsystem.

43. The method of claim 40, further comprising not updating the response indication in response to address packets requesting a shared access right to the coherency unit.

44. The method of claim 40, further comprising the memory subsystem sending a packet corresponding to the address packet to the interface if:

the response indication indicates that the memory should not respond;

the coherency unit is in a shared global access state in the node; and

the address packet requests write access to the coherency unit.

45. The method of claim 44, further comprising the interface sending a coherency message requesting write access to the coherency unit to the different node identified by the node identifier in response to receiving the packet corresponding to the address packet.

46. The method of claim 40, further comprising the memory subsystem sending a packet corresponding to the address packet to the interface if:

the response indication indicates that the memory should not respond; and

the coherency unit is in an invalid global access state in the node.

47. The method of claim 33, further comprising the node updating the node identifier in response to the interface receiving coherency messages from other nodes in the multi-node system.

48. The method of claim 33, further comprising the active device that sends the address packet specifying the coherency unit gaining the access right to the coherency unit in response to receiving the data packet from the memory subsystem if the memory subsystem sends the data packet.